

UNITED STATES PATENT AND TRADEMARK OFFICE

---

BEFORE THE BOARD OF PATENT APPEALS  
AND INTERFERENCES

---

*EX PARTE HOPEN ET AL.*

U.S. PATENT APPLICATION NUMBER 10/733,808

FILING DATE: DECEMBER 10, 2003

---

BRIEF ON APPEAL

---

CARR & FERRELL *LLP*  
2200 GENG ROAD  
PALO ALTO, CA 94303  
T: 650.812.3400  
F: 650.812.3444

ATTORNEY FOR THE APPELLANTS  
AND REAL-PARTY-IN-INTEREST

## TABLE OF CONTENTS

REAL-PARTY-IN-INTEREST .....	3
RELATED APPEALS AND INTERFERENCES .....	4
STATUS OF THE CLAIMS .....	4
STATUS OF AMENDMENTS .....	4
SUMMARY OF THE CLAIMED SUBJECT MATTER .....	5
GROUND OF REJECTION TO BE REVIEWED ON APPEAL .....	11
ARGUMENT .....	12
CONCLUSION AND REQUESTED RELIEF .....	16
CLAIMS APPENDIX .....	17
EVIDENCE APPENDIX .....	23
RELATED PROCEEDINGS APPENDIX .....	24

**REAL-PARTY-IN-INTEREST**  
**(37 C.F.R. § 41.37(c)(1)(i))**

The Appellants in the present appeal are Chris Hopen, Gary Tomlinson, John Brooke, Derek Brown , Jonathan Burdge, and Rodger Erickson—the joint inventors of U.S. patent application 10/733,808 (the '808 Application). The real-party-in-interest and assignee of record is Aventail LLC, a subsidiary of SonicWALL, Inc.

SonicWALL designs, develops, manufactures, and sells network security, content security, and business continuity solutions for businesses of various sizes.

SonicWALL's products are designed to provide secure Internet access to both wired and wireless broadband customers, enable secure Internet-based connectivity for distributed organizations, inspect the content entering and leaving its customers networks, protect organizations against inbound and outbound e-mail threats, and provide business continuity in the case of data or connectivity loss. SonicWALL products are designed to provide Internet security solutions for networks ranging in size from one to more than 15,000 users; enterprises having branch offices, telecommuting employees or point-of-sale (POS) locations, and e-commerce applications that handle millions of secure transactions daily.

An assignment by and between the Appellants and Aventail Corporation is recorded at Reel 015564 and Frame 0319. A merger agreement between Aventail Corporation and Aventail LLC is recorded at Reel 022200 and Frame 0475. Aventail LLC is the owner of the entire right, title, and interest in the '808 Application.

**RELATED APPEALS AND INTERFERENCES**  
**(37 C.F.R. § 41.37(c)(1)(ii))**

The '808 Application makes no claim for priority benefit. U.S. patent application numbers 11/903,219 filed September 21, 2007 (the '219 Application) and 11/927,321 filed October 29, 2007 (the '321 Application) claim the priority benefit of the '808 Application. The '219 Application is awaiting first examination. The '321 Application received a first office action on November 12, 2009. Neither the '219 Application or the '321 Application are presently involved in any appeal or interference proceeding before the Board of Patent Appeals and Interferences.

**STATUS OF THE CLAIMS**  
**(37 C.F.R. § 41.37(c)(1)(iii))**

Claims 1-12 and 14-23 are all pending and under appeal. Claims 13 and 24-25 have been cancelled. All of the aforementioned claims have been at least twice rejected. No claims have been allowed or are otherwise objected to by the Examiner.

**STATUS OF AMENDMENTS**  
**(37 C.F.R. § 41.37(c)(1)(iv))**

The amendments presented March 2, 2009 have been entered into the prosecution history; those amendments were presented subsequent to a non-final rejection dated October 1, 2008. A further response was filed July 27, 2009 and in response to a final rejection dated May 27, 2009. No amendment to the claims was presented in the July 27, 2009.

**SUMMARY OF THE CLAIMED SUBJECT MATTER**  
**(37 C.F.R. § 41.37(c)(1)(v))<sup>1</sup>**

Independent claim 1 as presented for appeal recites:

A network appliance, comprising:  
    at least one platform service stored in memory, the at least one platform server executable by a processing device;  
    a service monitor stored in memory and executable by a processor to monitor a working status of the at least one platform service when executed from memory, the service monitor using interprocess communications, the working status indicating that the at least one platform service is running, not running, or starting; and  
    a load balancer stored in memory and executable by a processor to perform load balancing on received communications based on at least the working status of the at least one platform service.

Independent claim 1 sets forth a network appliance that can provide a variety of software services including both platform services and load balancing services.

*Specification*, 6:18-28; 7:1-5. When using two or more network appliances, the network appliance described in the present application can replace a substantial portion of a convention network. *Specification*, 6:22-24.

The system of independent claim 1 is akin to Figure 4, which illustrates “a network appliance.” *Specification*, Figure 4; 12:4-12. Various embodiments of the network appliance “may be implemented using electronic hardware” or “by executing software instructions on a programmable computing device or computer.” *Specification*, 8:13-18. A computer system that can be used to implement the network appliance includes “a processing unit” and “a system memory.” *Specification*, Figure 2; 8:16-20.

---

<sup>1</sup> All references to the *SPECIFICATION* are exemplary and are not intended to be limiting. The present references are made solely to satisfy the requirements of 37 C.F.R. § 41.37(c)(1)(v). No reference is intended—nor should it be construed—as an admission or denial as to any requirement for patentability, including but not limited to those requirements set forth in 35 U.S.C. § 112, ¶ 1 as they pertain to written description and enablement.

“[T]he network appliance provides a load balancing service and one or more platform services.” *Specification*, 12:6-7. A load balancer “direct[s] incoming client communications to platform services based upon the status information provided by the service monitoring process, the interface monitoring process, and the node management process.” *Specification*, 19:14-28; 21:1-11. The network appliance also includes a service monitor that may check “the status of one or more of the platform services hosted by the network appliances” through “interprocess communications.” *Specification*, 22:26-27; 26:21-23. The service monitor gathers working status reports from a node management process that indicates the “status of each platform service as being “RUNNING,” “NOT RUNNING,” or “STARTING.”” *Specification*, 22:16-25.

Independent claim 11 as presented for appeal recites:

A system for balancing loads in a network, the system comprising:  
a first network appliance having:  
at least one first platform service stored in memory,  
a service monitor stored in memory and executable to monitor a working status of the at least one first platform service when executed from memory, the service monitor using interprocess communications, the working status indicating that the at least platform service is running, not running, or starting; and  
a first load balancer stored in memory and executable to perform load balancing on communications received by the first network appliance based on at least the working status of the at least one first platform service, the working status indicating whether the at least one platform service is running; and  
a second network appliance having:  
at least one second platform service stored in memory, and  
a second load balancer stored in memory and executable to perform load balancing on communications received by the second network appliance, wherein the first network appliance is configured to receive all client communications to the network unless the first load balancer fails, and the second network appliance is configured to receive all client communications to the network if the first load balancer fails.

Independent claim 11 sets forth a system for load balancing on a network where if “the load balancing service of a network appliance fails, another network appliance can provide load balancing. *Specification*, 6:27-28, 7:1.

The system of independent claim 11 is akin to Figure 4, which illustrates “a network appliance.” *Specification*, Figure 4; 12:4-12. Various embodiments of the network appliance “may be implemented using electronic hardware” or “by executing software instructions on a programmable computing device or computer.” *Specification*, 8:13-18. A computer system that can be used to implement the network appliance includes “a processing unit” and “a system memory.” *Specification*, Figure 2; 8:16-20.

“[T]he network appliance provides a load balancing service and one or more platform services.” *Specification*, 12:6-7. A load balancer “direct[s] incoming client communications to platform services based upon the status information provided by the service monitoring process, the interface monitoring process, and the node management process.” *Specification*, 19:14-28; 21:1-11. If “the load balancing service provided by the network appliance 307A is designated the primary load balancing service, then the load balancing service provided by the network appliance 307B will service as the secondary or “standby” load balancing service, and will operate only when the load balancing service of the network appliance 307A fails.” *Specification*, 17:11-20. If the “primary load balancing service fails, then the secondary or backup load balancing service can take over ownership ... and assume the responsibility of load balancing incoming communications.” *Specification*, 18:5-9. *Specification*, 17:11-20.

The network appliance also includes a service monitor that may check “the status of one or more of the platform services hosted by the network appliances” through “interprocess communications.” *Specification*, 22:26-27; 26:21-23. The service monitor gathers working status reports from a node management process that indicates the “status of each platform service as being “RUNNING,” “NOT RUNNING,” or “STARTING.”” *Specification*, 22:16-25.



Independent claim 20 recites:

A method of processing client communications to a network comprising:  
receiving a first client communication at a first network appliance hosting at least one first platform service;  
employing a load balancer hosted by the first network appliance to direct the first client communication to the at least one first platform service hosted by the first network appliance based on at least a working status of the at least one first platform service, the working status indicating that the at least one platform service is running, not running, or starting;  
receiving a second client communication at the first network appliance; and  
employing the load balancer to direct the second client communication to a second platform service hosted by a second network appliance based on at least the working status of the at least one first platform service and a working status of the second platform service.

Independent claim 20 sets forth a method for processing and routing client communications routed among multiple network appliances. *Specification*, 7:6-10.

The method of independent claim 20 relates to the “network appliance” illustrated in Figure 4. *Specification*, Figure 4; 12:4-12. “[T]he network appliance provides a load balancing service and one or more platform services.” *Specification*, 12:6-7. A load balancer “direct[s] incoming client communications to platform services based upon the status information provided by the service monitoring process, the interface monitoring process, and the node management process.” *Specification*, 19:14-28; 21:1-11. A load balancer acts upon a received communication and “directs it to the appropriate platform service.” *Specification*, 18:3-4. The platform services hosted by a first network appliance and a second network appliance “will concurrently process incoming communications from client computers.” *Specification*, 19:3-5. A load balancer will route new client communications to both the platform services provided by its own

network appliance and the platform services provided by the other network appliance.”  
*Specification*, 19:5-13.

The network appliance also includes a service monitor that may check “the status of one or more of the platform services hosted by the network appliances” through “interprocess communications.” *Specification*, 22:26-27; 26:21-23. The service monitor gathers working status reports from a node management process that indicates the “status of each platform service as being “RUNNING,” “NOT RUNNING,” or “STARTING.”” *Specification*, 22:16-25.

**GROUND OF REJECTION TO BE REVIEWED ON APPEAL**

**(37 C.F.R. § 41.37(C)(1)(vi))**

- I. Has the Examiner Evidenced the Presence of Each and Every Claimed Element in the Independent Claims and a Corresponding *Prima Facie* Case of Obviousness with Respect to the Combination of *O'Neil* and *Wilding*?
- II. Has the Examiner Established Proper Motivation to Combine and a Corresponding *Prima Facie* Case of Obviousness with Respect to the Combination of *O'Neil* and *Wilding*?

**ARGUMENT**  
**(37 C.F.R. § 41.37(c)(1)(vii))**

**I. THE EXAMINER HAS NOT EVIDENCED THE PRESENCE OF EACH AND EVERY CLAIMED ELEMENT IN THE INDEPENDENT CLAIMS AND A CORRESPONDING *PRIMA FACIE* CASE OF OBVIOUSNESS WITH RESPECT TO THE COMBINATION OF *O'NEIL* AND *WILDING***

The Examiner rejects independent claims 1, 11, and 20 under 35 U.S.C. § 103(a) as unpatentable over U.S. patent number 6,128,279 (hereinafter *O'Neil*) in view of U.S. patent application number 2003/0212788 (hereinafter *Wilding*). The Applicants respectfully traverse in that *O'Neil* and *Wilding* fail to disclose each and every element of the independent claims.

To support a conclusion that the claim would have been obvious requires that **all** the claimed elements were known in the prior art and that one skilled in the art could have combined those elements. See *KSR v. Teleflex*, 127 S.Ct. 1727, 1739 (2007); see also MPEP § 2143. *O'Neil* and *Wilding* – individually and collectively – fail to disclose at least ‘a load balancer stored in memory and executable by a processor to perform load balancing on received communications based on at least the working status of the at least one platform service,’ wherein ‘the working status’ refers to ‘running, not running, or starting.’

Ascertaining the differences between the prior art and the claims at issue requires considering both the invention and the prior art references **as a whole**. MPEP § 2141.02. When read as a whole, the claimed invention is not merely “monitoring a working status,” as posited by the Examiner. *May 27, 2009 Office Action*, 3, ¶ 8. The claims, when read as a whole, clearly require that ‘load balancing’ be ‘based on at least the working status’ of ‘running, not running, or starting.’

Neither *O'Neil* nor *Wilding* collectively disclose the Applicants' claimed 'load balancing' based on 'working status' of 'running, not running, or starting.' With respect to *O'Neil*, the Examiner expressly admits that *O'Neil* "does not explicitly teach of a working status indicating that the at least one platform service is running, not running, or starting." *May 27, 2009 Office Action*, 3, ¶ 7. With respect to *Wilding*, the Examiner notes instances of 'working status,' but fails to cite to any 'load balancing' based on 'working status' of 'running, not running, or starting.'

*Wilding* merely teaches a "generic control interface" that can "assess the service's operability, aliveness, and availability" "without requiring a detailed understanding of the specific operations necessary for controlling or monitoring the specific service." *Wilding*, Abstract. *Wilding* notes that "[i]n the event that a fault monitor detects an abnormality in the service that it is monitoring, the fault monitor takes corrective action." *Wilding* [0027]. Such corrective action may include "start the service," "send a notification to the system administrator to alert the administrator to the absence of the service," or "wait a short period of time and then re-evaluate the service to determine if the service has returned to an available status, failing which it may notify the system administrator." *Wilding* [0036]. None of the "corrective actions" of *Wilding* and that the Examiner appears to correlate to the Applicants' claimed 'working status' involve **any load balancing**; any such correlation with load balancing is, therefore, in error. Notwithstanding, there is no reason to combine *Wilding* with any aspect of 'load balancing' (regardless of whether it is disclosed in the art) other than to derive the Applicants' claimed invention, which is an impermissible use of hindsight analysis.

The Examiner does not dispute that "[n]either *O'Neil* nor *Wilding* teach of (sic) the claimed load balancing based on a working status indication that the platform service is running, not running, or starting." *Advisory Action*, 2. The Examiner argues, instead, that Applicant's argument for reading the claim language as a whole is "arguing against the references individually." *Advisory Action*, 2. The Applicant respectfully disagrees.

The Applicants' previously presented argument in addition to the aforementioned remarks merely indicates that the Applicants have considered the claimed invention in context and as a whole.

The Applicants respectfully submit that the Examiner's Section 103 arguments indicate a failure to consider the references **and the claimed invention** as a whole. See *May 27, 2009 Office Action*, 3, ¶ 7; see also *Advisory Action*, 2. As such, the rejection runs contrary to the MPEP's admonition that individual claim terms not be cherry-picked or read in isolation. See MPEP § 2141.02. Such distillation of claim language is improper and "disregards the statutory requirement that the invention be viewed 'as a whole.'" *Jones v. Hardy*, 727 F.2d 1524, 1530 (Fed. Cir. 1984).

When viewed and considered as a whole, the claimed invention clearly requires that 'load balancing' be based on 'working status' of 'running, not running, or starting,' which is not disclosed in *O'Neil* or *Wilding*, individually or in combination.

## **II. THE EXAMINER HAS NOT ESTABLISHED A PROPER MOTIVATION TO COMBINE AND A CORRESPONDING *PRIMA FACIE* CASE OF OBVIOUSNESS WITH RESPECT TO THE COMBINATION OF *O'NEIL* AND *WILDING***

Pursuant to the Supreme Court's decision in *KSR v. Teleflex*, the Examiner must still provide some basis for a proposed combination (*i.e.*, motivation) regardless of whether the elements were previously known to exist in the art. "[K]nowledge of a problem and motivation to solve it are entirely different from motivation to combine particular references to reach the particular claimed method." *Innogenetics, N.V. v. Abbot Labs.*, 512 F.3d 1363 (Fed. Cir. 2008) (affirming that "[a] generalized motivation to develop a method is not the kind of motivation required by the patent laws").

A mere conclusory statement that modifications of the prior art to meet the claimed invention would have been well within the ordinary skill of the art “because the references relied upon teach that all aspects of the claimed invention were individually known in the art is **not sufficient** to establish a *prima facie* case of obviousness without some **objective reason to combine** the teachings of the references.” *Ex parte Levengood*, 28 USPQ2d 1300 (Bd. Pat. App. & Inter. 1993) (emphasis added); see also MPEP § 2142.01(iv). “Combining prior art references without evidence of such a suggestion, teaching, or motivation simply takes the inventor’s disclosure as a blueprint for piecing together the prior art to defeat patentability—the essence of hindsight.” *In re Dembiczak*, 175 F.3d 994, 999 (Fed. Cir. 1999).

The Applicants respectfully submit that the Examiner has used impermissible hindsight to reach a conclusion of obviousness, since the Examiner has picked a reference (*i.e.*, *Wilding*) that has nothing to do with ‘load balancing’ in an (albeit failed) effort to disclose the ‘working status’ that is admittedly missing from *O’Neil*. A person of ordinary skill in the art looking to improve upon “load balancing” would not look to incorporate *Wilding*. Further, as *Wilding* does not have anything to do with ‘load balancing,’ the incorporation of such would change the operation and purpose of *Wilding*. See *Wilding*, Abstract.

## CONCLUSION AND REQUESTED RELIEF

Because both *O'Neil* and *Wilding* fail to teach the claimed 'load balancing' 'based on at least the working status' of 'running, not running, and starting,' *O'Neil* and *Wilding* therefore fails to teach each and every claim limitation of the independent claims. Any claim dependent upon the aforementioned independent claims —either directly or via an intermediate dependent claim—is allowable for at least the same reasons as the independent claim from which it depends. As such, each and every one of the dependent claims of the present application are also in condition for allowance. For at least these reasons, the Examiner's rejection should be withdrawn.

Respectfully submitted,  
Chris A. Hopen et al.

December 28, 2009

By:           /Tam Thanh Pham/            
Tam Thanh Pham (Reg. No. 50,565)  
**CARR & FERRELL LLP**  
2200 Geng Road  
Palo Alto, CA 94303  
T: 650.812.3400  
F: 650.812.3444

ATTORNEY FOR THE APPELLANT  
AND REAL-PARTY-IN-INTEREST



**CLAIMS APPENDIX**  
**(37 C.F.R. § 41.37(c)(1)(viii))**

1. A network appliance, comprising:
  - at least one platform service stored in memory, the at least one platform server executable by a processing device;
  - a service monitor stored in memory and executable by a processor to monitor a working status of the at least one platform service when executed from memory, the service monitor using interprocess communications, the working status indicating that the at least one platform service is running, not running, or starting; and
  - a load balancer stored in memory and executable by a processor to perform load balancing on received communications based on at least the working status of the at least one platform service.
2. The network appliance of claim 1, further comprising a backplane interface through which the network appliance exchanges data with another device.
3. The network appliance of claim 2, wherein the another device hosts at least one second platform service stored in memory, and the service monitor is executable to monitor a working status of the at least one second platform service using communications transmitted over the backplane.
4. The network appliance of claim 1, further comprising an interface monitor stored in memory and executable to monitor a working status of interfaces and connections employed by the network appliance.
5. The network appliance of claim 1, wherein the at least one platform service is an access method service.

6. The network appliance of claim 5, wherein the access method service is a virtual private network service.
7. The network appliance of claim 5, wherein the access method service is an extranet Web service.
8. The network appliance of claim 1, further comprising a node manager stored in memory and executable to determine the working status of the at least one platform service and provide the determined working status of the at least one platform service to the service monitor.
9. The network appliance of claim 1, further comprising a distributed cache service that caches information relating to at least one platform service on another network appliance.
10. The network appliance of claim 9, wherein the at least one platform service is an access method service, and the information cached includes authentication information and encryption key information for encryption sessions hosted by the access method service.

11. A system for balancing loads in a network, the system comprising:
- a first network appliance having:
    - at least one first platform service stored in memory,
    - a service monitor stored in memory and executable to monitor a working status of the at least one first platform service when executed from memory, the service monitor using interprocess communications, the working status indicating that the at least platform service is running, not running, or starting; and
    - a first load balancer stored in memory and executable to perform load balancing on communications received by the first network appliance based on at least the working status of the at least one first platform service, the working status indicating whether the at least one platform service is running; and
  - a second network appliance having:
    - at least one second platform service stored in memory, and
    - a second load balancer stored in memory and executable to perform load balancing on communications received by the second network appliance,
- wherein the first network appliance is configured to receive all client communications to the network unless the first load balancer fails, and the second network appliance is configured to receive all client communications to the network if the first load balancer fails.
12. The system of claim 11, wherein the second network appliance further includes a second service monitor stored in memory and executable to monitor a working status of the at least one second platform service using interprocess communications.
14. The system of claim 11, wherein the at least one first platform service is an access method service.

15. The system of claim 14, wherein the access method service is a virtual private network service.

16. The system of claim 14, wherein the access method service is an extranet Web service.

17. The system of claim 11, wherein the at least one second platform service is an access method service.

18. The system of claim 17, wherein the access method service is a virtual private network service.

19. The system of claim 17, wherein the access method service is an extranet Web service.

20. A method of processing client communications to a network comprising:

- receiving a first client communication at a first network appliance hosting at least one first platform service;

- employing a load balancer hosted by the first network appliance to direct the first client communication to the at least one first platform service hosted by the first network appliance based on at least a working status of the at least one first platform service, the working status indicating that the at least one platform service is running, not running, or starting;

- receiving a second client communication at the first network appliance; and

- employing the load balancer to direct the second client communication to a second platform service hosted by a second network appliance based on at least the working status of the at least one first platform service and a working status of the second platform service.

21. The method of claim 20, further comprising:

- analyzing the first client communication to determine if the first client communication includes association data indicating that the first client communication is associated with the at least one first platform service; and

- determining that the first client communication includes association data indicating that the first communication is associated with the at least one first platform service.

22. The method of claim 21, wherein the association data is a session identifier

identifying an encryption session maintained by the at least one first platform service.

23. The method of claim 20, further comprising:

executing a load balancing algorithm to determine whether the second client communication should be directed to the second platform service; and

determining that the second client communication should be directed to the second platform service based upon results of the executed load balancing algorithm.

**EVIDENCE APPENDIX**  
**37 C.F.R. § 41.37(c)(1)(ix)**

Not applicable in the present appeal.

**RELATED PROCEEDINGS APPENDIX**

**37 C.F.R. § 41.37(c)(1)(x)**

Not applicable in the present appeal.